Qualitative Spatial Reasoning: Extracting and Reasoning with Spatial Aggregates

Chris Bailey-Kellogg
Dept. of Computer Sciences, Purdue University
250 N. University St.
West Lafayette, IN 47907, USA
cbk@cs.purdue.edu

Feng Zhao
Palo Alto Research Center (PARC)
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
fz@alum.mit.edu

Abstract

Reasoning about spatial data is a key task in many applications, including geographic information systems, meteorological and fluid flow analysis, computer-aided design, and protein structure databases. Such applications often require the identification and manipulation of *qualitative* spatial representations, for example, to detect whether one "object" will soon occlude another in a digital image, or to efficiently determine relationships between a proposed road and wetland regions in a geographic data set. Qualitative spatial reasoning (QSR) provides representational primitives (a spatial "vocabulary") and inference mechanisms for these tasks. This paper first reviews representative work on QSR for *data-poor* scenarios, where the goal is to *design* representations that can answer qualitative queries without much numerical information. It then turns to the *data-rich* case, where the goal is to *derive* and manipulate qualitative spatial representations that efficiently and correctly abstract important spatial aspects of the underlying data, for use in subsequent tasks. This paper focuses on how a particular QSR system, Spatial Aggregation (SA), can help answer spatial queries for scientific and engineering data sets. A case study application of weather analysis illustrates the effective representation and reasoning supported by both data-poor and data-rich forms of QSR.

Introduction

The ability to perceive spatial objects and reason about their relations seems effortless for humans, but has proved so difficult for computers that they have yet to attain the capabilities of a 5-year-old child. Part of the computational problem lies in the difficulty of *identifying and manipulating qualitative spatial representations*. For example, while the pixels in a digital image implicitly define the locations of spatial "objects," the task at hand might require a more qualitative characterization of the configuration of these objects, say, whether one object will soon occlude another. Handling spatial data is a key task in many applications, including geographic information systems (GIS), meteorological and fluid flow analysis, CAD systems, and protein structure databases (Fig. 1 (Zhao *et al.* 1999)). For example, a GIS system might have large amounts of numerical information about spatial features such as highways and terrain, but require query mechanisms to efficiently determine qualitative relationships such as those between a proposed route and wetlands regions.

Qualitative spatial reasoning (QSR) addresses these problems with representational primitives (a spatial "vocabulary") and inference mechanisms. QSR approaches can be characterized by two important and complementary classes of problems. Problems in the first class are data poor, and the goal is to *design* representations that can answer qualitative queries without much numerical information. This addresses an important aspect of common sense reasoning by humans, and can be found in many practical applications such as computer-aided tutoring or diagram understanding. Because of lack of detailed numerical information, representations used by the approaches to data-poor problems are often carefully designed by hand with respect to a task at hand. Problems in the second class (e.g. scientific and engineering applications from fluid flow analysis to distributed control design) are data rich, and the goal is to *derive* and manipulate qualitative spatial representations that efficiently and correctly abstract important spatial aspects of the underlying data, and that can be used for subsequent tasks. The approaches to data-rich problems are complementary to those for data-poor problems in that they can *automatically construct* spatial representations. Computational efficiency in reasoning arises from appropriate qualitative spatial representations; for example, a qualitative description of a temperature distribution as a configuration of iso-contours focuses the search for good thermal control designs. Similarly, qualitative representations allow efficient access and manipulation of data, for example, correlating maps (e.g. a road map, a utilities map, and a forestry map) in a GIS system, determining interaction of parts in a CAD design, and planning paths for a robotics application. An important feature of qualitative spatial representations is their ability to relate reasoning results to underlying data (rich or poor) and domain knowledge provided by the user.

In this paper, we will first review the representative work on QSR for data-poor scenarios. We will then turn to the data-rich case, and focus on how a particular QSR system, Spatial Aggregation (SA), can help answer spatial queries for scientific and engineering data sets. Finally, we will present a particular application that illustrates the effective representation and reasoning supported by both forms of QSR.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **2003**	2. REPORT TYPE	3. DATES COVERED **00-00-2003 to 00-00-2003**
4. TITLE AND SUBTITLE **Qualitative Spatial Reasoning: Extracting and Reasoning with Spatial Aggregates**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Purdue University,Dept. of Computer Sciences,250 N. University St,West Lafayette,IN,47907**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT **see report**		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**12**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

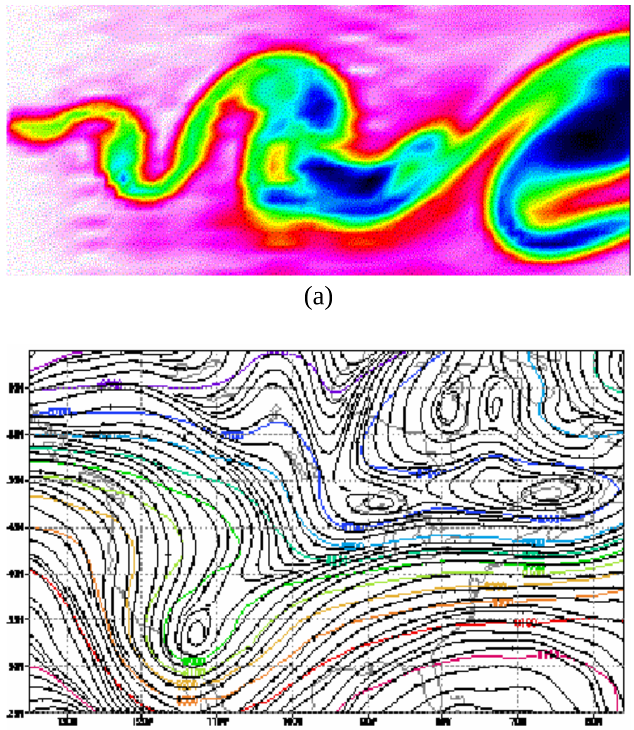

(a)

(b)

Figure 1: Spatial objects in physical fields: (a) A fluid flow. The fluid field describes how objects such as high density regions and large vortex structures are spatially distributed (shown here) and temporally evolving (not shown). (b) A 300mb weather map over North America. The data in a typical meteorological map includes pressure, temperature, and wind velocities on a spatial grid. An experienced meteorologist could identify qualitatively important weather features such as the location of a cold front and the direction of its movement, by extracting and correlating geometric features such as pressure troughs and thermal packing.

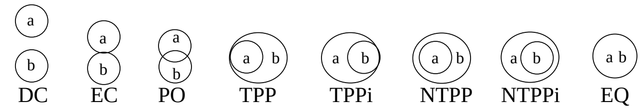

Figure 2: The region-connection calculus (RCC) represents and manipulates pairwise relationships between regions of space.

QSR for Data-Poor Problems

Qualitative reasoning (QR) research uses high-level representations of physical systems and domain knowledge for tasks such as prediction, diagnosis, reconfiguration, and tutoring (Forbus 1984; de Kleer & Brown 1984; Kuipers 1986; Weld & de Kleer 1990; Falkenhainer & Forbus 1991), without requiring significant amounts or quality of data. Classical QR work deals primarily with temporal aspects of a system, abstracting away its spatial properties. Following the spirit of this QR research, work in QSR for data-poor domains has focused on similarly "minimalist" spatial representations and inference mechanisms.

Much QSR work has studied purely topological descriptions of spatial regions and their relationships. These approaches often seek to generalize Allen's temporal interval calculus (e.g. a before b, a overlaps b, etc.) (Allen 1983) into higher-dimensional, spatial relationships. One representative approach, the Region-Connection Calculus (RCC) (Cui, Cohn, & Randell 1992), provides predicates for expressing and reasoning about the relationships among topological regions (arbitrarily-shaped chunks of space). One version, the RCC-8, provides eight jointly exhaustive and pairwise disjoint predicates (Fig. 2): disconnected from (DC), externally connected to (EC), partially overlaps (PO), tangential proper part of (TPP and TPPi), nontangential proper part of (NTPP and NTPPi), and identical with (EQ). The axioms specifying these relationships provide rigorous underpinnings to support spatial reasoning. For example, boolean functions (e.g. union, intersection, and difference) allow composition of complex spatial objects (i.e. topological shapes). Additional predicates can then test, via theorem proving, topological features of these objects (e.g. connected, number of holes) and feasibility of additional relationships. An important characteristic of the predicates is that they support reasoning about *continuity* — a temporal process between two regions must pass through the possible relationships in a well-defined way (e.g. they can't go directly from being disconnected to identical). While the most general RCC theories are undecidable, tractable subsets suitable for various domains have been identified and applied to applications ranging from GIS to visual programming languages.

In addition to topological relationships, QSR researchers have also studied other key qualitative aspects of spatial objects, such as size and shape, and relationships, such as orientation and distance. A full review of these representations is beyond the scope of this paper (see e.g. (Cohn & Hazarika 2001)). Shape representations typically go beyond pure topology to specify some amount of geometric information, such as convex/concave portions of a boundary, and

use multi-scale representations similar to those described in the next section. Uncertainty in shape can be handled with coupled RCC-like predicates specifying an object's certain interior and uncertain exterior. Distance, orientation, and size can be represented relatively, for example, indicating order-of-magnitude relationships and rules for combining them (e.g. distances sum when oriented properly). Of particular interest is that these representations must often bring in some amount of metric information in order to make significant inference possible.

In fact, Forbus, Nielsen, and Faltings (Forbus, Nielsen, & Faltings 1991) hypothesized that some metric information is necessary in general for qualitative spatial reasoning. More precisely, the *poverty conjecture* states that "there is no problem-independent, purely qualitative representation of space or shape." Purely qualitative means essentially that no detailed metric information supporting perceptual-like processing is available. Problem independence means that the representation must be general — a small set of spatial objects constrained to only specific interactions in a specific domain might indeed admit a purely qualitative representation, but that representation might then break down with the addition or modification of a single part.

In order to balance the tension between qualitativeness and generality, the Metric Diagram / Place Vocabulary (MD/PV) theory (Forbus, Nielsen, & Faltings 1991) takes a hybrid approach, linking together metric information supporting quantitative queries, with sets of special qualitatively-important entities (places) in a domain. For example, in analysis of clock mechanisms (Forbus, Nielsen, & Faltings 1991), the place vocabulary is computed for pairs of interacting parts, which are specified with CAD-like metric diagrams that can determine such interactions. This approach addresses a key concern in qualitative reasoning, ensuring the appropriateness of a choice of qualitative vocabulary, since the place vocabulary is computed for a specific problem. It also ensures tight coupling between qualitative and quantitative aspects of the reasoning. In an approach that is similar at a high level, although different in application details, the Spatial Semantic Hierarchy (Kuipers & Levitt 1988; Kuipers 2000) discovers "interesting" locations in the construction of mappings between topological and metric maps for robot navigation.

Qualitative physical fields (Lundell 1996) capture spatially distributed qualitative parameters, that is, each spatial region consists of a uniformly-valued qualitative feature. For example, a model of heating might describe regions as being warm or cold, as well as regions that are sunny or shaded. Note that, as opposed to pairwise interactions, this representation is at its very heart continuous. This representation supports reasoning about spatio-temporal processes, in an extension of Qualitative Process Theory (Forbus 1984). Continuing the example, a qualitative heat flow would be established between topologically adjacent regions of different qualitative temperature (from warm to cold). This heat flow would establish a temporal process changing the "front" between the two regions and ultimately the associated temperatures. This approach will be discussed in more detail in the case study section.

QSR for Data-Rich Problems

In contrast with the data-poor application domains discussed above, many important science and engineering applications are characterized by massive amounts of spatially-distributed numerical data (Fig. 1). For example, in order to predict the weather, meteorologists utilize pressure, temperature, and wind velocity data collected from a large number of spatially distributed weather sensors. Similarly, in designing aircraft with minimal drag, engineers study wind tunnel and simulation data specifying airflow over a body at many points and over many instants of time. In these applications, geometric as well as topological characterizations are necessary; for example, a temperature field is influenced by the geometry of the domain, spatial variations in material property, and boundary conditions.

The massive amount of data, either collected from experiments or produced by simulations, poses significant computational challenges that can be addressed by QSR. In particular, a *central problem* is the automatic construction of qualitative spatial representations from a given data set, in order to focus the search space for data interpretation and design tasks. QSR approaches to these problems are often built upon a sound mathematical theory of geometric and topological analysis, for example, the theory of cell complexes from algebraic topology that naturally defines "closeness", "composition", and "abstraction" (Munkres 1984). As discussed earlier, the availability of data provides us the opportunity to automate the construction of qualitative spatial representations. QSR differs from traditional numerical methods for spatial data analysis problems that also abstract numerical data at multiple levels of resolution. For example, engineers use multigrid methods (Briggs 1987) to analyze numerical properties of physical phenomena using a hierarchy of grid discretization. The main difference between QSR and numerical methods lies in the ontological abstraction QSR adopts. QSR supports more abstract, qualitative reasoning by introducing notions of objects that explicitly encapsulate key spatial properties of a physical domain. For example, meteorologists use abstract structures such as isobars, pressure troughs, and pressure cells to reason about the underlying pressure data at a higher level of abstraction. This key insight — *physical properties such as continuity and locality give rise to regions of uniformity in spatially distributed data* — enables QSR to overcome the challenge of massive data. In fact, this insight is similar to that underlying the MD/PV approach described in the previous section. Domain-specific physical knowledge justifies the extraction of qualitative information in support of more abstract reasoning processes.

QSR in data-rich domains has many connections and parallels to work in scientific visualization (Rosenblum & others 1994) and scientific data mining (Ramakrishnan & Grama 2001). For example, weather data can be visualized using pseudocolor to represent temperature, iso-contours to connect points of equal pressure, needle diagrams to indicate directions of wind flow with arrows, streamlines to show connected flows, and animations of these to show changes over time. Interactive visualizations allow scientists to explore, focus, filter, project, and transform large data sets.

Feature detection algorithms (e.g. for vortices in fluid data) both identify and track spatial structures over time (Samtaney *et al.* 1994; Yip 1995; Junker & Braunschweig 1995; Ordóñez & Zhao 2000). Similarly, in scientific data mining, algorithms seek to cluster, generalize, and classify patterns and correlations in databases. For example, mining such patterns can allow identification of general climate patterns across regions (Lu, Han, & Ooi 1993), automatic cataloging of sky images (Fayyad, Weir, & Djorgovski 1993), recognition of volcanoes in images of the surface of Venus (Burl *et al.* 1994), and tracking of cyclones in weather data (Stolorz & others 1995).

Spatial Aggregation

A particular QSR approach for data-massive domains, called Spatial Aggregation (SA) (Yip & Zhao 1996) follows an imagistic reasoning (Yip, Zhao, & Sacks 1995) style, applying vision-like routines to manipulate multi-layer geometric and topological structures in spatially distributed data. This allows it to leverage the connections described above with visualization and data mining. In the spirit of qualitative reasoning, however, it focuses on explicit representation and manipulation of objects, explainability of results, and utilization of explicitly encoded domain knowledge. Spatial Aggregation is partially motivated by some of the spatial reasoning problems raised by Abelson *et al.* (Abelson *et al.* 1989). The Abelson paper describes a number of approaches to interpreting numerical results of simulations of dynamical systems. These problems often possess a set of geometric and topological constraints that can be exploited to significantly cut down the search space, and can be used to communicate the interpretation results to human experts. For example, in interpreting the qualitative behaviors of a nonlinear dynamical system, one can describe the set of trajectories that share the same asymptotic behaviors as a flow pipe, a geometric object that can be easily visualized (Yip & Zhao 1996). Several early examples of using geometric and spatial reasoning to aid in scientific computation include KAM (Yip 1991), which interprets the behaviors of Hamiltonian systems, MAPS (Zhao 1994), which designs control laws based on a geometric analysis of the state equations of a dynamical system, and HIPAIR (Joskowicz & Sacks 1991), which analyzes the kinematics of fixed-axis mechanisms.

SA organizes computation around image-like representations of spatially distributed data (Fig. 1). In the *field ontology*, the input is a *field* mapping from one continuum to another. For example, a 2-D temperature field associates a temperature with each point, mapping from $\mathbb{R}^2$ to $\mathbb{R}^1$; a 2-D fluid field associates a velocity with each point, mapping from $\mathbb{R}^2$ to $\mathbb{R}^2$. A field is information-rich, in that its representation requires many bits. The identification of structures in a field (e.g. iso-bars, pressure cells, and fronts) is a form of data reduction: the data-rich field representation is abstracted into a more concise structural representation. For example, a set of points on a curve can be described more compactly by a parameterized spline — the spline parameters are a much more concise representation than the enumeration of points. Note that the qualitative physical field approach described in the previous section starts with an abstract description of a field (qualitative domain and range in the field); the synergy between the data-rich and data-poor fields will be explored further in the case study section.

SA uncovers structures in fields at multiple levels of abstraction, with the structures uncovered at one level becoming the input to the structure-discovery process at the next level. For example, in a weather data analysis application (Huang & Zhao 2000), SA could extract from pressure data the isobars, pressure cells, and pressure troughs. As discussed above, continuities in a field give rise to regions of uniformity that can be abstracted as spatial structures (e.g. isothermal contours are connected curves of equal (or similar enough) temperature). Similarly, these structures exhibit their own continuities; therefore, multi-layer structures arise from continuities in fields at multiple scales. Spatial objects are introduced as primitives in QSR to encapsulate the geometric and topological properties of these points, curves, regions, or volumes. Mathematically, a spatial object is a *cell* — a portion of space topologically equivalent to a ball (Munkres 1984). Adjacency between the objects is defined by the contiguity of their cells. Navigating the mapping from field to abstract description through multiple layers rather than in one giant step allows the construction of modular programs with manageable pieces that can use similar processing techniques at different levels of abstraction. The multi-level mapping also allows higher-level layers to use global properties of lower-level objects as local properties of the higher-level objects. For example, the average temperature in a region is a global property when considered with respect to the temperature data points, but a local property when considered with respect to a more abstract region description.

SA provides a set of data types and operators for constructing the spatial aggregate hierarchy. The data types and operators make explicit use of domain-specific knowledge (see Fig. 3), in particular similarity and closeness of both field objects and their features which are encoded with metrics, adjacency relations, and equivalence predicates. Yip and Zhao (Yip & Zhao 1996) present a number of application programs, ranging from dynamical systems analysis to mechanical mechanism analysis, in terms of the same set of generic operators parameterized by different such domain knowledge. The central data type of SA, the *neighborhood graph*, is an explicit representation of an object adjacency relation. The definition of adjacency is domain-specific and depends on the metric properties of the input field. Common adjacency relations include Delaunay triangulations, minimal spanning trees, and uniform grids. The neighborhood graph serves as computational glue, localizing interactions between neighboring objects. The main SA operators *aggregate* objects into neighborhood graphs satisfying an adjacency predicate, *classify* neighboring nodes into equivalences classes with respect to an equivalence predicate specifying domain-specific feature similarity, and *redescribe* equivalence classes into higher-level objects with respect to a domain-specific abstraction mechanism. Additional operators search through neighborhood graphs, check consistency of objects, extract geometric properties, and so forth. By instantiating these operators with proper knowl-

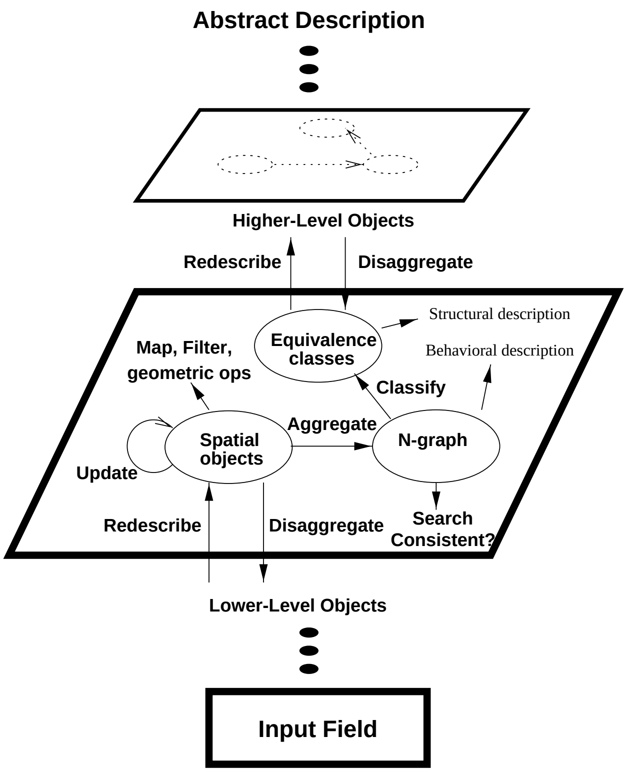

Figure 3: Spatial Aggregation: objects are created, transformed, and abstracted using a set of generic operators.

edge at different levels of abstraction, SA allows specification of a variety of application programs.

The Spatial Aggregation Language (SAL) (Bailey-Kellogg, Zhao, & Yip 1996; Yip & Zhao 1996), summarized in Table 1, implements the SA theory with a C++ library of data types and operators and an interpreted, interactive environment layered over the library. The library supports construction of efficient C++ programs, while the interpreter supports rapid programming of modeling tasks by providing a convenient, high-level interface. SAL lets programmers conveniently explore trade-offs in the specification of domain knowledge such as neighborhood relations and equivalence predicates, and interactively and graphically examine and modify the results. For those interested, the source code of SAL can be downloaded from `www.cs.purdue.edu/homes/cbk/sal.html` or `www.parc.com/zhao/sal.html`.

Recent work on SA (Bailey-Kellogg & Ramakrishnan 2001; Ramakrishnan & Bailey-Kellogg 2002) has moved to bridge the quantitative-qualitative gap in the opposite direction, using qualitative structures to guide sample selection in the underlying field. This approach of sparse data mining is particularly important for applications where very expensive data collection must be carefully planned (e.g. for fluid dynamics simulation and aircraft design). In particular, the iterative approach performs spatial analysis of data in one iteration, identifies ambiguities arising in the anal-

- *Primitive Objects* represent locations and structures in spatial data.

  Example: 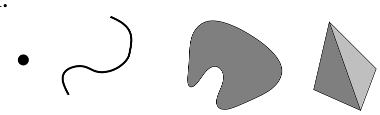

- *Compound Objects* combine primitive objects.
  - *Spaces* group objects.

    Example (points and curves): 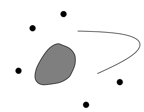
  - *Fields* associate objects and features.

    Example (point $\mapsto$ temperature): 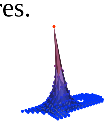
  - *Ngraphs* relate nearby objects.

    Example (Delaunay triangulation): 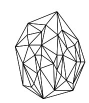
  - *Equivalence classes* group similar objects.

    Example (by vector direction): 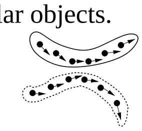

- *Means of Abstraction* connect compound objects at one level of abstraction and primitive objects at the next.

  Example (by convex hull): 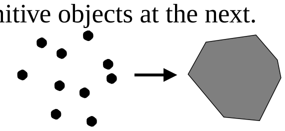

Table 1: Components of the Spatial Aggregation Language.

ysis, and focuses sample selection in the next iteration so as to clarify the ambiguities, maximizing information content and improving the analysis. This approach has been shown to make highly effective, explainable sampling decisions in several case studies, including discovery of “pockets” in n-dimensional space by aggregation of gradient vector fields in an interpolated representation derived from a minimal number of targeted samples (Bailey-Kellogg & Ramakrishnan 2001); analysis of matrices via a perturbation sampling approach (Ramakrishnan & Bailey-Kellogg 2002), using consistent correspondence of features to determine properties such as Jordan form from a small number of samples; influence-based model decomposition for decentralized control design (Bailey-Kellogg & Zhao 1998; 1999; 2001; Bailey-Kellogg & Ramakrishnan 2001), where locality of a few sampled control effects supports high-quality decomposition of problem domains and reasoning about tradeoffs among computation, communication among decentralized controls, and resulting control quality.

Case Study: Reasoning with Weather Data

Consider the approach taken by meteorologists interpreting weather data in order to make predictions about future conditions. They make sense of large, multi-category data sets by recognizing and explicitly labeling aggregate weather features such as high/low pressure centers, pressure troughs, thermal packings, fronts, and jet streams (Fig. 4 (Huang & Zhao 2000)). They then use weather rules, such as the following, in order to correlate these features and establish prediction patterns:

1. “Major and minor 500mb troughs are good indicators of existing or potential adverse weather.” (Air Weather Service 1975)
2. “At 850mb, the polar front is located parallel to and on the warm side of the thermal packing.” (Air Weather Service 1975)
3. “A front lies in a pressure trough and the iso-bars make an abrupt change in direction at the front.” (Blair & Fite 1965)
4. “A front moves slowly when it is nearly parallel to the iso-bars and increases in velocity as the number of iso-bars intersecting it increases.” (Blair & Fite 1965)
5. “A strong high east of a low, especially if the high is increasing in intensity or is nearly stationary, will retard the low or deflect it to the left or right. Two lows close together tend to unite.” (Blair & Fite 1965)

These rules are intuitively expressed in terms of animated, interacting objects that have rich spatial and physical properties and often defy concise mathematical characterization.

This weather data analysis application illustrates a more general class of practical modeling and design problems requiring an understanding of distributed parameter fields. The task of modeling such fields is important by itself, for example in ecology applications, where researchers desire to understand interactions among different parameters (e.g. shade and temperature profile) (Lundell 1996). As discussed above, modeling also leads directly to prediction, based on

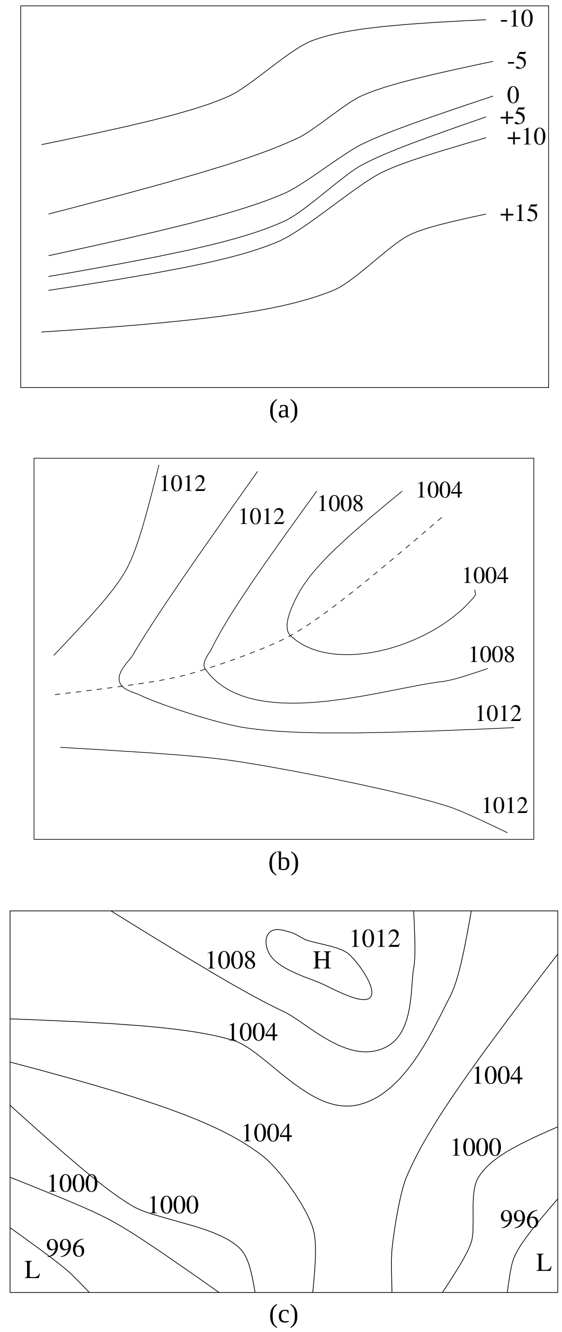

Figure 4: Weather features: (a) a thermal packing; (b) a pressure trough, where solid lines representing iso-bars and dashed line trough position; (c) highs “H” and lows “L” in a pressure map.

features extracted from a field at one point in time, or even better the history of such features over time (Yip 1997; Ordóñez & Zhao 2000). While the weather data application isn't amenable to control (yet!), modeling is also important when engineers desire to regulate a physical field with some set of controls (Bailey-Kellogg & Zhao 2001), for example guiding a robotic laser welding arm in response to temperature data from an infrared camera (Doumanidis 1997) or maintaining a uniform temperature profile with a set of concentric circles of heat lamps (Kailath *et al.* 1996). In the following two subsections, we discuss how data-poor QSR supports reasoning with models of such heat flows, and how data-rich QSR supports extraction and manipulation of such models.

Data-Poor Reasoning

This section follows the approach taken by Lundell for qualitative physical fields (Lundell 1995; 1996). In reasoning about temperature fields in ecology, dense, precise numerical data and corresponding models are often not available. However, it is desirable to envision qualitative differences in the temporal evolution for a given (qualitative) model. As described above, the model is defined in terms of a field — an association of some parameter (here temperature) with spatial objects in the domain. Composite fields are defined by spatial interactions of fields with overlapping domains and different parameters (e.g. temperature and pressure). The static definition of a field's domain and involved parameters is specified separately from the dynamic process capturing the interactions among the parameters over time.

The value of a field parameter belongs to some qualitative space of possible values (e.g. "warm" or "cold"). At a point in time, the field domain is partitioned into maximal regions of the same value (e.g. iso-thermal regions). This is essentially a place vocabulary (see the discussion above of MD/PV). The resulting qualitative spatial representations are then amenable to the general QSR techniques described in the preceding section. In particular, composite fields (e.g. temperature and shade) are naturally computable via intersection of regions in the separate fields. Fig. 5 illustrates such qualitative physical fields for the temperature modeling application, in both the underlying geometric domain and a diagrammatic representation that captures the relevant topological connectivity and continuity.

Dynamic changes in fields (e.g. due to heat flow) are captured with a spatio-temporal extension to the Qualitative Process Theory (Forbus 1984). In particular, spatio-temporal processes are captured as interactions between spatial and temporal processes. In the case study application, each region is warmed by a temporal process that takes into account their irradiation and temperature differences, and the regions of boundaries are adjusted by a spatial process that considers the differences in temperatures. More precisely, the temporal process manipulates variables for heating rate in the regions, with a negative qualitative relationship between the heating rate and the amount of shade (i.e. more shade means less heating from the sun) and a monotonic influence between the temperature in the region and the heating rate. The spatial process then distributes heat

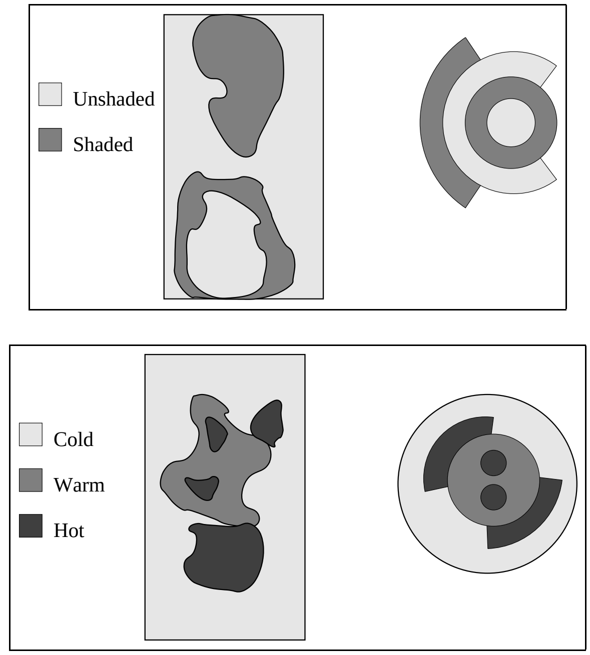

Figure 5: Example qualitative physical fields.

via flow between adjacent regions of different temperatures. It is specified in terms of an "expansion region" of applicability, starting at the boundary between two regions and spreading at a rate proportional to the temperature difference until equilibrium is reached. Temporal processes within the expansion region cool one part and heat the other. Simulation of this set of processes yields an envisionment of the qualitatively interesting state transitions in the evolution of a temperature/shade field.

Data-Rich Modeling

This section illustrates that qualitative reasoning about physical fields can extract rich structure from large spatial datasets, in support of tasks such as prediction and analysis. In particular, we focus on the identification of troughs and ridges in weather data. We provide here a high-level discussion of a Spatial Aggregation-based approach; interested readers are referred to (Huang & Zhao 2000) for details about the approach and results.

Troughs and ridges are important features in weather analysis; for example, high-altitude troughs give rise to the bending of jet streams and are important for extended weather forecasts, while surface troughs are useful for locating weather fronts. Trough features are only qualitatively understood (see the sample weather prediction rules at the start of this section); sometimes experts even give different answers about the existence of a trough in a weather map. The key to identifying troughs lies in the qualitative structures of a field of atmospheric pressure data. In particular, the shape and configuration of *iso-bars* — the iso-contours of the pressure field, collecting points of the same pressure value — indicate areas where troughs are likely present. Visually, troughs and ridges are stacks of iso-bar

segments bending sharply and consistently to one direction, with troughs pointing away from lower iso-bars and ridges away from higher iso-bars. Fig. 4(b) shows a trough. Due to the Coriolis force, winds tend to follow iso-bars, so sharply bending iso-bars indicate sharp change of wind direction, which usually causes more advection, more mixing of warm air with cold air, and therefore deteriorating weather.

As previously discussed, data-rich QSR focuses on the extraction and manipulation of structures in spatial data. These structures arise as equivalence classes of neighboring objects according to some similarity measure, redescribed as primitive objects at a higher level of abstraction for further analysis. A Spatial Aggregation-based trough-finding algorithm uses this approach to extract the same qualitative spatial features that experts do — sharply-bending segments of iso-curves of pressure data. The input to the algorithm is a gridded pressure dataset, and the output is a contoured pressure chart with troughs labeled. In a preprocessing step, iso-bar points are interpolated from the gridded data, yielding a set of "iso-points," with pressure at specified contour levels. The algorithm then proceeds through two levels of aggregation (see again the section on Spatial Aggregation for a description of operations in a typical level), the first to group points into iso-bars and the second to group segments of iso-bars into troughs and ridges. We briefly describe the Spatial Aggregation steps in this process (using slightly different wording from the original paper); Fig. 6 (Huang & Zhao 2000) illustrates with data from a 500mb pressure data set from a National Weather Service data server.

Level I:

1. *Aggregate*: Build a Delaunay triangulation neighborhood graph for the iso-points. The Delaunay triangulation has a number of important geometric properties. Most importantly here, while the triangulation only acts on iso-points, its edges are "well-behaved" with respect to the aggregated iso-curves, in that its edges connect only points within a single curve or in two topologically-adjacent curves.
2. *Classify*: Form equivalence classes of neighboring points sharing the same pressure value. At the same time, classify graph edges into "strong" adjacencies, connecting same-class points, and "weak" adjacencies, connecting points in different classes.
3. *Redescribe*: Abstract each class of same-pressure points into an iso-curve object. As previously discussed, forming a higher-level object allows the computation of aggregate properties; here, curvature is an especially important property. We note that while other algorithms (e.g. marching cubes (Lorenson & Cline 1987)) can also be used to contour a pressure dataset, the approach taken here yields more structure in the spatial objects, and this structure will prove useful in later steps.

Level II:

1. *Filter*: Segment the curves and extract high-curvature curve segments. Curve segmentation breaks a curve into piece-wise simple parameterized curves (e.g. straight lines or circular arcs); for example, a split-and-merge algorithm (Pavlidis & Horowitz 1974) splits curves at places with high approximation errors and merges segments to avoid over-segmentation. Simple thresholding then allows extraction of high-curvature segments.
2. *Aggregate*: Build a neighborhood graph for the high-curvature curve segments based on the between-class ("weak") adjacencies from Level I. That is, two curves are adjacent if and only if a constituent point in one is adjacent to a constituent point in the other in the original Delaunay triangulation.
3. *Classify*: Form equivalence classes of neighboring curve segments that bend in similar directions, within some tolerance.
4. *Redescribe*: Abstract equivalence classes of similar-direction high-curvature curve segments into troughs. The abstraction process constructs a curve (e.g. B-spline) through the stack of iso-curves (e.g. through a representative point on each).

As Fig. 7 (Huang & Zhao 2000) demonstrates, the results of this algorithm are in qualitative agreement with those of professional meteorologists. While the expert-drawn trough seems smoother and more visually pleasing, the exact shape and position are not as important for a synoptic map at this scale.

This trough-finding algorithm illustrates the importance of explicitly representing and reasoning with multi-level spatial structures. The edges in a neighborhood graph play two distinct roles in the aggregation process. The classification process in Level I above uses domain-specific knowledge to distinguish these roles as *strong* and *weak* adjacencies (Fig. 8 (Huang & Zhao 2000)). Strong adjacencies carry information about the interactions and connections *among the constituent parts* of an aggregate spatial object, and are abstracted as structural information of the object. Weak adjacencies carry information about the interactions and connections *between* aggregate objects, and are abstracted into a higher-level neighborhood graph. Explicitly representing these adjacencies allows the programmer to use a natural encoding of domain knowledge, in terms of equivalence predicates, in order to identify objects that are internally connected, externally bounded, and related at multiple levels of abstraction.

Conclusions and Future Research Directions

We have described several representative approaches to data-poor and data-rich problems in QSR. In many applications dealing with spatial data, qualitative spatial representations and inferences are preferable because either detailed numerical information is not available for the domain or existing numerical methods are unable to describe the kinds of geometric and topological structures in data sets that can help answer high-level spatial queries. As the sample applications demonstrate, QSR is an important aspect of common sense reasoning and can have a significant impact on many technical and scientific applications.

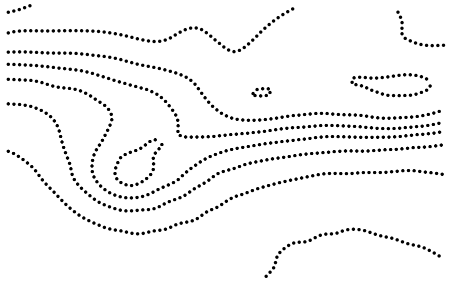

(a) *Preprocessed input*: iso-points

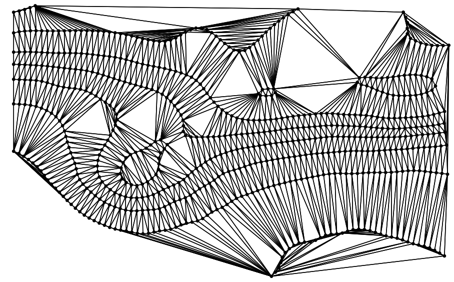

(b) *Aggregate* (I): point triangulation

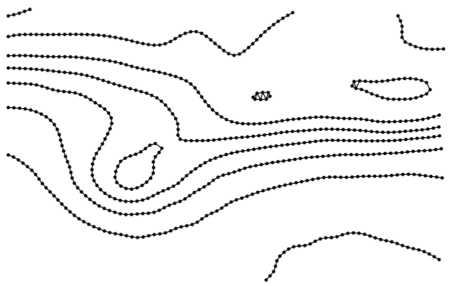

(c) *Classify* (I): pressure classes

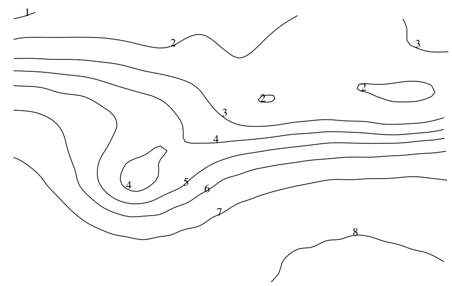

(d) *Redescribe* (I$\mapsto$II): iso-bars

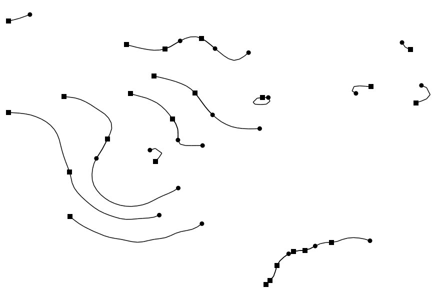

(e) *Filter* (II): high-curvature segments

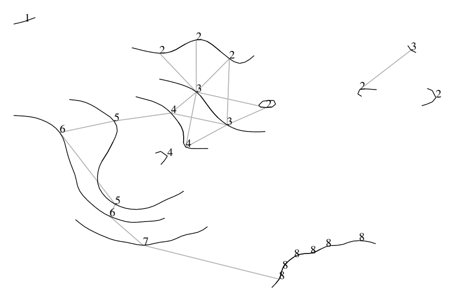

(f) *Aggregate* (II): constituent point adjacency

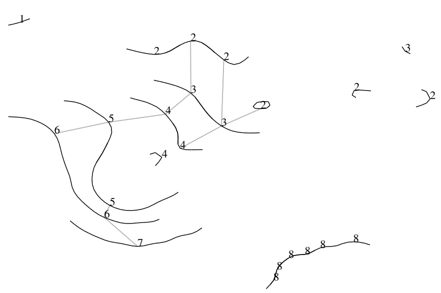

(g) *Classify* (II): similar curvature

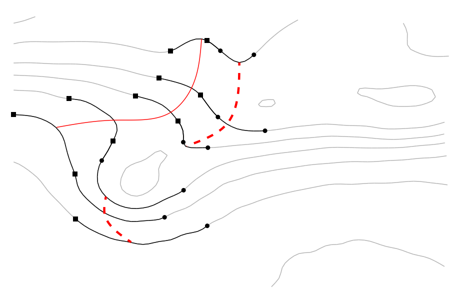

(h) *Redescribe* (II): troughs and ridge
(troughs: dashed; ridge: solid)

Figure 6: Example of a Spatial Aggregation approach to extraction of pressure troughs and ridges.

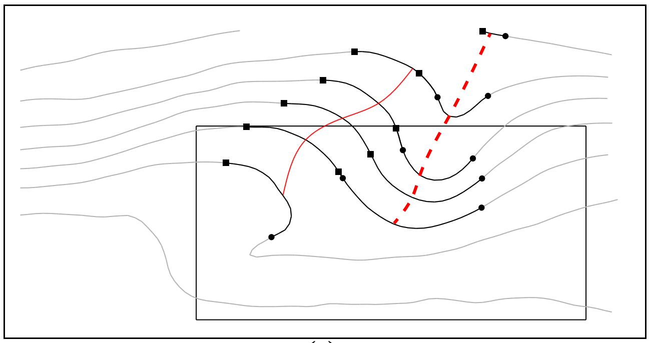

(a)

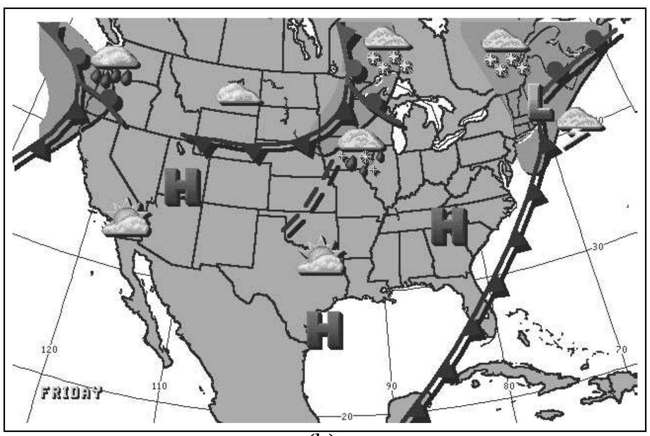

(b)

Figure 7: Labeling a weather chart: (a) The high-altitude trough (dashed line) detected by the Spatial Aggregation-based algorithm. (b) The corresponding trough (parallel dashed line) drawn by meteorologists for the national weather forecast map for roughly the same area as the small box in (a).

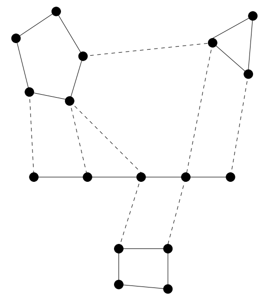

(a)

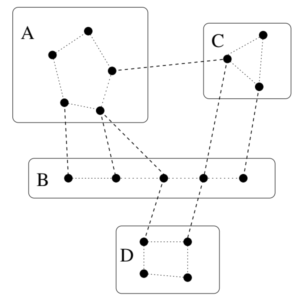

(b)

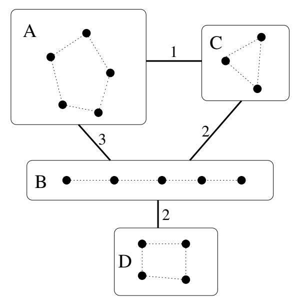

(c)

Figure 8: Explict representation and aggregation of spatial adjacencies. (a) Neighborhood graph edges are classified into strong (solid) and weak (dashed) adjacencies. (b) Strong adjacencies connect internal structure of aggregate objects A, B, and C. (c) A higher-level neighborhood graph is constructed from lower-level weak adjacencies between constituent objects.

QSR is a rich problem domain for QR research. To fully realize the potential of QSR, we will need to address a number of open research issues. For example, Spatial Aggregation introduces a number of spatial primitives for describing structures in a data-rich physical field. What are other formulations of the problem using say a primal-dual space representation? What are additional primitives and inference operators in SA that might be appropriate? How can probabilistic information be incorporated? An important problem in synthesizing the approaches to data-poor and data-rich problems is to use data-rich approaches to automatically build models for data-poor approaches. Here, the data used to build a model for a data-poor problem could perhaps come from a domain where the physics constraints are similar enough and numerical information is readily available. For example, in an ecology application, although detailed numerical information for a particular region might not be available due to lack of instrumentation for that region, the model building process could leverage data from another similar region where sensors have already been deployed.

While each of the approaches we have described is to some degree based on mathematical theories of topology and geometry, we have yet to develop a rigorous and formal basis for a general theory of qualitative spatial reasoning that can unify the different approaches. Equally important is the development of a set of problem characterizations that can aid in transforming a general theory into an efficient algorithm for a task.

Acknowledgement

Fig. 1, from (Zhao *et al.* 1999), is included courtesy of Ohmsha, Ltd. Figs. 4, 6, 7, and 8, from (Huang & Zhao 2000), are included courtesy of IOS Press. The Spatial Aggregation research described here was conducted at Ohio State and Palo Alto Research Center (PARC) by the authors and Xingang Huang and Iván Ordóñez. It was supported by NSF grant CCR-9457802, ONR grant N00014-97-1-0599, and a Sloan Research Fellowship.

References

Abelson, H.; Eisenberg, M.; Halfant, M.; Katzenelson, J.; Sacks, E.; Sussman, G.; Wisdom, J.; and Yip, K. 1989. Intelligence in scientific computing. *CACM* 32(5):546–562.

Air Weather Service. 1975. Back to basics. In *AWS FOT Seminar STT-Q9-0004*.

Allen, J. 1983. Maintaining knowledge about temporal intervals. *CACM* 26(11):832–843.

Bailey-Kellogg, C., and Ramakrishnan, N. 2001. Ambiguity-directed sampling for qualitative analysis of sparse data from spatially-distributed physical systems. In *The 17th Int'l Joint Conf. on Artificial Intelligence (IJCAI)*, 43–50.

Bailey-Kellogg, C., and Zhao, F. 1998. Qualitative analysis of distributed physical systems with applications to control synthesis. In *The 15th Nat'l Conf. on Artificial Intelligence (AAAI)*, 232–239.

Bailey-Kellogg, C., and Zhao, F. 1999. Influence-based model decomposition. In *The 16th Nat'l Conf. on Artificial Intelligence (AAAI)*, 402–409.

Bailey-Kellogg, C., and Zhao, F. 2001. Influence-based model decomposition. *Artificial Intelligence* 130(2):125–166.

Bailey-Kellogg, C.; Zhao, F.; and Yip, K. 1996. Spatial aggregation: language and applications. In *The 13th Nat'l Conf. on Artificial Intelligence (AAAI)*, 517–522.

Blair, T., and Fite, R. 1965. *Weather Elements*. Prentice-Hall.

Briggs, W. 1987. *A Multigrid Tutorial*. Lancaster Press.

Burl, M.; Fayyad, U.; Perona, P.; Smyth, P.; and Burl, M. 1994. Automating the hunt for volcanoes on Venus. In *Proceedings of Computer Vision and Pattern Recognition Conference (CVPR-94)*, 302–309.

Cohn, A., and Hazarika, S. 2001. Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae* 46(1-2):1–29.

Cui, Z.; Cohn, A.; and Randell, D. 1992. Qualitative simulation based on a logical formalism of space and time. In *The 10th Nat'l Conf. on Artificial Intelligence (AAAI)*, 679–684.

de Kleer, J., and Brown, J. 1984. A qualitative physics based on confluences. *Artificial Intelligence* 24(1–3):7–83.

Doumanidis, C. 1997. In-process control in thermal rapid prototyping. *IEEE Control Systems* 17(4):46–54.

Falkenhainer, B., and Forbus, K. 1991. Compositional modeling: Finding the right model for the job. *Artificial Intelligence* 51(1–3):95–143.

Fayyad, U.; Weir, N.; and Djorgovski, S. 1993. SKICAT: a machine learning system for the automated cataloging of large-scale sky surveys. In *Proceedings of 10th International Conference on Machine Learning*, 112–119.

Forbus, K.; Nielsen, P.; and Faltings, B. 1991. Qualitative spatial reasoning: the CLOCK project. *Artificial Intelligence* 51(1–3):417–471.

Forbus, K. 1984. Qualitative process theory. *Artificial Intelligence* 24(1–3):85–168.

Huang, X., and Zhao, F. 2000. Relation-Based Aggregation: Finding Objects in Large Spatial D atasets. *Intelligent Data Analysis* 4(2):129–147.

Joskowicz, L., and Sacks, E. 1991. Computational kinematics. *Artificial Intelligence* 51(1–3):381–416.

Junker, U., and Braunschweig, B. 1995. History-based interpretation of finite element simulations of seismic wave fields. In *Proceedings of IJCAI*, 1789–1797.

Kailath, T.; Schaper, C.; Cho, Y.; Gyugyi, P.; Norman, S.; Park, P.; Boyd, S.; Franklin, G.; Saraswat, K.; Moslehi, M.; and Davis, C. 1996. Control for advanced semiconductor device manufacturing: A case history. In Levine, W., ed., *The Control Handbook*, 1243–1259. CRC Press.

Kuipers, B., and Levitt, T. 1988. Navigation and mapping in large-scale space. *AI Magazine* 9(2):25–43.

Kuipers, B. 1986. Qualitative simulation. *Artificial Intelligence* 29(3):289–338.

Kuipers, B. 2000. The spatial semantic hierarchy. *Artificial Intelligence* 119(1–2):191–233.

Lorenson, W., and Cline, H. 1987. Marching cubes: a high resolution 3d surface construction algorithm. *Computer Graphics* 21(4):163–169.

Lu, W.; Han, J.; and Ooi, B. 1993. Discovery of general knowledge in large spatial databases. In *Proceedings of Far East Workshop on Geographic Information Systems*, 275–289.

Lundell, M. 1995. A qualitative model of gradient flow in a spatially distributed parameter. In *Proceedings of International Qualitative Reasoning Workshop*, 131–138.

Lundell, M. 1996. A qualitative model of physical fields. In *Proceedings of AAAI*, 1016–1021.

Munkres, J. 1984. *Elements of Algebraic Topology*. Addison-Wesley.

Ordóñez, I., and Zhao, F. 2000. STA: Spatio-Temporal Aggregation with Applications to Analysis of Diffusion-Reaction Phenomena. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI'00)*, 517–523.

Pavlidis, T., and Horowitz, S. 1974. Segmentation of plane curves. *IEEE Transactions on Computers* C-22(8):860–870.

Ramakrishnan, N., and Bailey-Kellogg, C. 2002. Sampling strategies for mining in data-scarce domains. *IEEE Computing in Science & Engineering (CiSE)* 4(4):31–43.

Ramakrishnan, N., and Grama, A. 2001. Mining Scientific Data. *Advances in Computers* Vol. 55:pages 119–169.

Rosenblum, L., et al., eds. 1994. *Scientific Visualization: Advances and Challenges*. Academic Press.

Samtaney, R.; Silver, D.; Zabusky, N.; and Cao, J. 1994. Visualizing features and tracking their evolution. *IEEE Computer Magazine* 27(7):20–27.

Stolorz, P., et al. 1995. Fast spatio-temporal data mining of large geophysical datasets. In *Proceedings of First International Conference on Knowledge Discovery and Data Mining*, 300–305.

Weld, D., and de Kleer, J., eds. 1990. *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann.

Yip, K., and Zhao, F. 1996. Spatial aggregation: theory and applications. *Journal of Artificial Intelligence Research* 5:1–26.

Yip, K.; Zhao, F.; and Sacks, E. 1995. Imagistic reasoning. *ACM Computing Surveys* 27(3):363–365.

Yip, K. 1991. *KAM: A system for intelligently guiding numerical experimentation by computer*. MIT Press.

Yip, K. 1995. Reasoning about fluid motion: finding structures. In *Proceedings of IJCAI*, 1782–1788.

Yip, K. 1997. Structural inferences from massive datasets. In *Proceedings of IJCAI*, 534–541.

Zhao, F.; Bailey-Kellogg, C.; Huang, X.; and I. Ordóñez. 1999. Intelligent simulation tools for mining large scientific data s ets. *New Generation Computing* 17(4):333–347.

Zhao, F. 1994. Extracting and representing qualitative behaviors of complex systems in phase spaces. *Artificial Intelligence* 69(1-2):51–92.

Author Information

Chris Bailey-Kellogg

Chris Bailey-Kellogg is an assistant professor of computer science at Purdue. He earned a BS/MS in electrical engineering and computer science from MIT and a PhD in computer and information science with Feng Zhao at Ohio State and Xerox PARC. He conducted postdoctoral research in computational biology with Bruce Donald at Dartmouth. His research combines geometric, symbolic, and probabilistic approaches for data analysis and experiment planning in scientific and engineering domains. He received an NSF CAREER award in 2003.

Feng Zhao

Feng Zhao is a Principal Scientist at Palo Alto Research Center (PARC), where he directs the Embedded Collaborative Computing Area (`www.parc.com/ecc`) in the Systems and Practices Laboratory. He is also Consulting Associate Professor of Computer Science at Stanford. His research interest includes sensor networks, diagnostics, qualitative reasoning, and control of dynamical systems.

Dr. Zhao received his PhD in Electrical Engineering and Computer Science from MIT in 1992. From 1992 to 1999, he was Assistant and Associate Professor of Computer and Information Science at Ohio State University. He received the ONR Young Investigator Award and the NSF Young Investigator Award, and was an Alfred P. Sloan Research Fellow in Computer Science. He serves on editorial boards of AI Magazine, IEEE Transaction on Signal Processing, IEEE Transaction on Control Systems Technology, New Generation Computing, and co-chaired the 1st and 2nd International Workshops on Information Processing in Sensor Networks (IPSN: `www.parc.com/events/ipsn03`). He has authored or co-authored over 70 peer-reviewed technical papers and is a co-inventor of three US Patents and five pending patent applications.